

N 9 3 - 2 7 0 4 7

**WHITE PAPER
on
SCIENCE OPERATIONS**

incorporating inputs from the
Science Operations MOWG

Ethan J. Schreier
Space Telescope Science Institute

April 1990

SCIENCE OPERATIONS WHITE PAPER

1. Introduction
2. Background
3. Recommendations in Data Analysis and Computing
 - 3.1 Data Access
 - 3.2 Software Support
 - 3.3 Workstation Support and Access
 - 3.4 Supercomputing Support and Access
 - 3.5 Network Support, Access and Coordination
 - 3.6 Electronic Publishing — Data, Abstracts
4. Recommendations in Operations
5. Management and Policy
6. Plan for Specific NASA- Related Activities
 - 6.1 General
 - 6.2 Data Analysis Support
 - 6.3 Operations Concept Studies
 - 6.4 Operations Technology Prototyping

1. INTRODUCTION

Major changes are taking place in the way astronomy gets done. There are continuing advances in observational capabilities across the frequency spectrum, involving both ground-based and space-based facilities. There is also very rapid evolution of relevant computing and data management technologies. However, although the new technologies are filtering in to the astronomy community, and astronomers are looking at their computing needs in new ways, there is little coordination or coherent policy. Furthermore, although there is great awareness of the evolving technologies in the arena of operations, much of the existing operations infrastructure is ill-suited to take advantage of them. Astronomy, especially space astronomy, has often been at the cutting edge of computer use in data reduction and image analysis, but has been somewhat removed from advanced applications in operations, which have tended to be implemented by industry rather than by the end user scientists.

It is likely that technology developments will continue to take place far more rapidly than most individual astronomers and new facilities will be able to take advantage of them, and increased attention to this problem is necessary. The challenge will be to provide new methodologies and infrastructures commensurate with the new technologies. It is likely that the impact which technological developments will have on astronomy over the next decade will be dominated by management and policy issues, and not by the technology itself. Furthermore, current procurement policies within the government introduce considerable time lags between the introduction of new technologies and their availability to the astronomy community. These policies, which were more appropriate to the era of infrequent main-frame procurements, are not suitable for the rapidly evolving world of personal workstations and mini-supercomputers. They must be reviewed and modified.

The purpose of this paper is threefold. First, we briefly review the background and general status of astronomy-related computing in Section 2. Second, we make recommendations in 3 areas: we summarize recommendations in the areas of data analysis in Section 3; operations in Section 4 (directed primarily to NASA-related activities); and issues of management and policy in Section 5, believing that these must be addressed to enable technological progress and to proceed through the next decade. Finally, we recommend specific NASA-related work as part of the Astrotech-21 plans, to enable better science operations in the operations of the Great Observatories and in the lunar outpost era.

2. BACKGROUND

Traditionally, astronomers (with exceptions) tended to be somewhat behind the other physical science disciplines in utilizing state-of-the-art computing technology. This situation has changed greatly over the last two decades, due largely to the fact that the type of astronomical problems being studied began demanding observations and analysis of data at many wavelengths, decreasing the separation between different sub-disciplines of astronomy. Thus, radio and x-ray astronomers, originally coming into the field with physics backgrounds and more familiar with computing technology, have become better integrated into "main-stream" astronomy. Optical astronomers have started using x-ray and radio facilities as well as data from the "near-optical" space facilities such as IUE and IRAS. Optical astronomy itself has had to face the task of processing high-volume digital data from CCDs and preparing in a major way for space astronomy with HST. The paucity of new space astronomy missions has made utilization of archival data more desirable. Theoretical modelers have gained access to supercomputers which allowed them to generate meaningful simulations and compare them with observational data in the various wavelength bands.

The process of cross-fertilization between the different astronomy disciplines has been accelerated by the computer networks which very recently began to make serious inroads into the astronomy community, even if underfunded and often uncoordinated by the primary astronomy funding agencies and centers. Astronomers thus became exposed, often reluctantly and inefficiently, to a variety of computing environments and they have been forced to think about more powerful and cost-effective types of computing hardware, data storage, networking, and even improved software development methodologies. It is thus not surprising that a growing number of astronomers started looking at computing and data management problems in a broader way, recognizing the existence of common problems, the possibility of common solutions, and above all, the need for better coordination and more funding.

The developing new outlook on data systems and computation in the astronomy community over the last decade has been manifested in several ways. First, individual astronomers and groups with enough support began to experiment with the newly introduced hardware, including workstations, PCs, mini-supercomputers, and supercomputers. They discovered ways to connect with existing networks. They started studying the use of optical disks for archiving large amounts of data. Sharing of software became more common, as the larger groups began to move away from "one-of-a-kind" solutions to computation problems, and by the early 1980s, the second-generation user-oriented data analysis systems (e.g. AIPS, IRAF, MIDAS) were conceptualized and/or under development. However, these were more-or-less grass-roots efforts carried out at the major user facilities, and there was only minimal funding and little effort toward cross-disciplinary coordination.

In recognition of the lack of adequate attention to the situation, the Committee On Data Management and Computation (CODMAC) was established in the late 1970s under the auspices of the Space Science Board. In several reports CODMAC assessed the magnitude of the computing and data problems facing the space science community, and made numerous observations and recommendations concerning the advantages of portable software, remotely accessible archives and wide-area networking, pointing the way toward distributed data analysis systems. But foremost among the findings were the conclusions that the problems standing in the way of qualitative improvements were mostly with management, not computing technology, and that user involvement in all stages of data system development was key to the achievement of usable capabilities. Although CODMAC findings were often given lip-service, for several years there was little true management attention to these latter problems.

A major positive step was taken by NASA in 1987 when the Astrophysics Division convened an Astrophysics Data System Study. This study, chaired by Gayle Squibb, and incorporating wide community participation, issued a report containing numerous recommendations for both specific activities and general guidelines to be followed, similar to those of CODMAC. Also included was a suggested architecture for an overall astronomy data system. The Astrophysics Division has started implementing several of the recommendations, via a dedicated Science Operations Branch, and with community interaction via a Science Operations Management Operations Working Group (SOMOWG). Key among the actions have been the establishment of an Astrophysics Data System pilot project and a peer-reviewed Software and Research Aids Program to support community efforts in astrophysics-related computing technologies. Increased attention has also been given to network links between different astronomy sites.

On the NSF side, there has been little new activity in astronomy-specific computing. Although some of the most substantial work toward community-wide data analysis systems was started at NSF-funded national centers (the FITS data interchange standard and the AIPS development at NRAO, and the IRAF development at NOAO), the redirection of NSF funding away from astronomy has limited these efforts, and, in fact, NASA is now helping to subsidize IRAF maintenance and development, and the FITS data standards. On the other hand, NSF's supercomputer centers and the related networking efforts have been very beneficial to astronomers. In fact, between the new NSF links, the existing ARPAnet, and NASA's SPAN and TCP/IP connections, there has been a dynamic qualitative improvement in connectivity in the astronomy community over the last few years.

Many individual astronomers and groups not only have taken advantage of the new capabilities but have recognized the advantages of moving toward general shared facilities. When the STScI was deliberating on a data analysis system for HST, it decided to build its data analysis software within the IRAF environment developed at NOAO, recognizing the advantages to the astronomer of not having to learn an unending stream of new data analysis systems. Similarly, when the ROSAT project also decided to use IRAF, it extended this philosophy to what had previously been an entirely separate sub-discipline of astronomy. The relevant groups have chosen to coordinate further developments with NASA's endorsement. In a related development, when HST needed to develop an optical disk archive capability, a facility was designed which could be used by multiple institutions, and was, in fact, developed by the STScI with substantial support by the ST-ECF at ESO and the Dominion Astrophysical Observatory in Victoria.

Despite the improvements over the last decade, much work remains to be done. The continuously evolving new technology must continue to be exploited, not only in small projects and in science computing, but more difficultly, in major facilities and in operations systems. It has become clear that large programs have a qualitatively harder job adopting new technologies and the new methodologies needed to exploit them, due most probably to the increased separation between the users and developers of systems as well as the usual management problems of large projects. This is particularly true in operations, which are handled on a mission by mission basis, and where for the most part there has been little grass roots effort to apply the lessons being learned in the data processing and analysis area. It is equally clear that improvement is still needed in the same areas identified by CODMAC a decade ago, including primarily user involvement in all stages of the development of data systems.

3. RECOMMENDATIONS IN DATA ANALYSIS AND COMPUTING

3.1 Data Access

- There should be easy remote access to digital data located at distributed data centers. Centers should provide documentation and expertise in the use of the data, software, and databases via active researchers at these centers.
- There should be access to processed data and software tools from not only the great observatories, but also from other space-based missions and ground-based surveys. This will involve new policies for ground-based observatories.
- There should be recognition of the need for data analysis to take place at astronomers' home institutions, with electronic or physical distribution of data as appropriate, and direct receipt of data from active missions a possibility
- There should be support for maintenance of catalogs and databases, and the necessary software and expertise.
- Proprietary rights policies and related international agreements should be reviewed, with the goal of making data more rapidly and widely available to a broader community including amateurs, educators, etc.
- There is great benefit to be derived from adequate descriptive material as well as the data itself. This material should include definitions, descriptions of processing, etc. which often get lost in the archiving process but which are essential to the proper scientific use of the data; this is an essential element of the ADS concept.
- Users should be protected from being forced to learn a multitude of user interfaces. We should promote a philosophy which includes minimizing the number of independent analysis systems and encouraging software portability, on-line help, standard command structures, etc.

3.2 Software Support

- There should be ongoing software maintenance support.
- Calibration software should be portable and included in the analysis systems.
- There should be support of advanced software developments and expert systems for data analysis. These should emphasize utility to a broad multi-disciplinary community and include cross-mission capabilities.

3.3 Workstation Support and Access

- The broad need for workstations as part of a distributed computing environment should be acknowledged
- The obsolescence factor should be recognized. Project plans should include replacement of equipment in a finite amount of time

3.4 Supercomputing Support and Access

- There is need for access/coordination between NASA users and NSF centers.
- There should be support for the development of supercomputer algorithms and other advanced computing strategies in image processing and data analysis
- Mini-supercomputers should be made available at major user facilities.

3.5 Network Support, Access and Coordination

- NASA should take a more active role in Internet coordination across agencies.
- NASA should take a more active role in connecting data centers to the Internet and to the ADS activity.
- There should be better connection/coordination of the science networks with operations activities in NASA.

3.6 Electronic Publishing -- Data, Abstracts

- There should be a means for making "published" data computer accessible.
- Abstracts should be made available on-line.
- Electronic proposal submission and perhaps review via e-mail should be encouraged.

4. RECOMMENDATIONS IN OPERATIONS

- NASA should support the development of portable, distributed, user-friendly, transparent observation planning tools. These should be consistent with the telescience concepts of remote mission planning and operations.
- There should be adequate bandwidth and minimal communications restrictions for remote observations and data communications, be it on the ground, in orbit, or from a lunar base. The concept of "INTERNET to the moon" should be encouraged.
- NASA must modernize its mission operations and communications infrastructure, including distributed operations concepts and direct reception of data.
- There should be more attention to Operations within the Science Operations Branch.
- There must be better coordination of operations development with instrument and spacecraft h/w development. There should be direct, frequent interaction between end-users, designers, developers and managers in the implementation of new operations capabilities.

- There should be user involvement in the development of a second generation TDRSS or its equivalent for non-low earth orbit missions.
- The potential for direct operation of small missions or experiments should be studied. This would allow more efficient interaction between the user and the facility, and could reduce costs.

5. MANAGEMENT AND POLICY

Management and policy problems far outweigh technical issues.

- NASA Operations Infrastructure should be made compatible with a distributed service-oriented operations concept.
- There should be increased emphasis on small, rapid-turnaround inexpensive missions. Benefits include lower launch costs and more continuity in research programs.
- Science goals can be better met by taking a bottom-line approach to what is truly needed, and not artificially linking the science missions to other NASA goals.

e.g. tying astrophysics to the manned program, with its safety and communications overheads, is regarded as very deleterious to science.

e.g. TDRSS support to science is inadequate, given the low priority relative to manned and DOD missions.

- The end-user must be involved in all phases of project development, and there should be more accountability within NASA in developing science missions.
- The procurement procedures used in large NASA missions is often incompatible with attaining the desired science goals:
 - existing talent in the astronomy and advanced technology communities should be utilized and not neglected.
 - the extended procurement cycle for many systems (e.g. computers) which almost assures obsolescence on delivery must be changed.
 - development contractors are not responsible for long-term operability and maintainability and do not adequately plan for these parts of the life cycle.
 - software development methodology must change to better involve end-users in all stages. Rapid prototyping must replace conventional adversarial development schemes.
- Multiwavelength capabilities should be encouraged in NASA programs, both in instrument complement (e.g. via addition of monitors) and in operations concepts.
- NASA should become more involved in astronomy education, and should encourage active involvement of amateur astronomers in NASA astronomy programs. Libraries should be modernized.

6. PLAN FOR SPECIFIC NASA-RELATED ACTIVITIES

6.1 General

NASA has recently formulated the Astrotech 21 Program, with the general goal of developing the technology base for a "new century astronomy program", including astrophysics missions of the 21st Century, and with a specific goal of preparing for lunar-based astronomy. In the sections below, we discuss specific activities which NASA should carry out in the context of Astrotech 21. However, these

same technologies would also be relevant for non-lunar sitings, where remote and/or unattended operations and communications constraints exist. The requirements of a lunar outpost for these science operations capabilities will be stronger, but not unique. This should be regarded as an advantage, since the capabilities can thus be usefully prototyped in advance, in realistic but perhaps less extreme applications. For astrophysics operations, there will be unique opportunities to maximize the utility of the proposed prototyping by making use of existing and planned astrophysics missions as testbeds for lunar outpost concepts. This is consistent with the need to set up the data and science operations infrastructure for supporting lunar-based telescopes *before* sending telescopes to the moon. The infrastructure to be developed must include capabilities for automated mission planning and scheduling, autonomous monitoring of both science and engineering data, including dynamic command management and autonomous response (both protective and for unique scientific actions), and intelligent data compression and distribution mechanisms.

There have been very strong recommendations that the operations capabilities be developed with continuing and direct user involvement, since many of the capabilities will be specific to astrophysics applications. It will thus be clearly advantageous to carry out prototyping within OSSA in applications which are closely related to candidate lunar outpost missions. In particular, many of the concepts could be tried out in the context of the currently planned Great Observatories.

6.2 Data Analysis Support

The traditional data analysis and data processing model for NASA space missions has been a series of "levels" of processing that gradually homogenize the data and remove artifacts of the source of the data. In a rough sense, the "Level 0" process cleans up the data with regard to data drop-outs, formatting, compression, and timing related to the transmission from the satellite to the ground. "Level 1" processing encompasses the reorganization of data, application of calibrations, and routine algorithms such as attitude corrections to place the data in "scientifically useful" form. These functions, Level 0 and Level 1, are usually done at the data capture facility and/or within mission data centers. The "Level 2" processing is the scientific data analysis which is done by the scientist, often with the assistance of the mission which provides the algorithms, software, facilities for this activity.

The lines that distinguish these levels of data processing and analysis are becoming blurred at present, particularly between "Level 1" and "Level 2". On the one hand, increases in the computer power available to individual scientists and institutions have reduced the need for centralized processing of data. On the other hand, the mix of users and the levels of sophistication and/or familiarity with the data requires that services such as standard processing of data continue to be available. The trend, then, must be to allow a broader spectrum of user services. This will require missions to provide not only standard processed data, but also data processing and analysis tools in portable and interoperable forms and unprocessed data to those who desire the less digested form of the data.

The need to develop interface specifications to allow such portability and interoperability must become a high priority of the scientific community so that current and future missions will be able to function in a highly integrated environment. This is true for data, catalogs, databases, and for software. The framework for supporting the prototyping of advanced software and astronomy information handling techniques already exists within the Astrophysics Division's Software and Research Aids Program and in the Astrophysics Data System. These programs should be encouraged and augmented.

6.3 Operations Concept Studies

Studies of operating modes and requirements, including scheduling, command generation, coordination with discipline facilities, etc. are needed. These will include the study of existing technologies for applicability to generic astrophysics-related requirements, and participation in the candidate mission studies. It is important to have visibility into the missions which are being considered, to make sure that operations-oriented considerations are included, and to feed requirements back into the operations prototyping activities discussed below. Thus, the specific missions being planned in the post-Great-Observatories era, and especially for the proposed lunar outpost, should be reviewed and studied to better define the required operations capabilities.

6.4 Operations Technology Prototyping

6.4.1 Data Compression (both noiseless, with 100% retrievability and no compromise in accuracy, and with dynamically determined accuracy).

Data compression is the process of encoding ("compressing") a body of data into a smaller body of data. It must be possible for the compressed data to be decoded back to the original data or some acceptable approximation of the original data. Data compression is a tool that data system designers can use in overcoming communications bandwidth and storage limitations which would otherwise make it impossible or difficult to satisfy science users needs. The use of this tool must be traded off against added computational loads for decompression, increased data link Bit Error Rate performance, and the potential of losing some scientific information.

In the case of future astronomy missions, data compression might be used within the flight segment (orbiting or lunar-based) to acquire more data than the data buffers or space-ground link would be able to accommodate, either temporally averaged or for high rate data bursts. Large volumes of image data on the ground must be stored, archived and browsed. Data compression can be a tool in reducing media costs, and can in some cases enable electronic transfer of science data for primary distribution or for interactive browsing.

Remote operations of instruments, such as may be required for a lunar observatory, will likely involve visual aids. These visual aids will present imagery to "tele-operators," providing feedback on the status and configuration of equipment as well as the general health and state of the observational data and communications links. Such purposes often do not require high fidelity reproduction at the receiving end -- only that the image "look" like the original.

Data compression techniques can be described as being either lossless, in which case the original data can be fully recovered, or lossy, in which case data (although perhaps no information) are lost. Lossless compression techniques will typically produce compression ratios of 2:1 or 3:1. In most cases, lossless techniques will be used for the transmission or storage of the science data. Lossy techniques can produce compression ratios as high as 1:100, and with significant processing overhead can even go as high as 1:1000. These techniques are more appropriately applied to transmission of browse products and to visual feedback aids for remote operations. However, even for the actual science data, it is possible that astronomers will be faced with a trade-off between the use of a lossy data compression scheme and no data at all. In that case, the lossy scheme is obviously the choice, provided that it is information-preserving with respect to the scientific purpose.

Significant progress has been made in the past several decades in the development of data compression algorithms and implementation of those algorithms in high speed hardware. In fact many internationally recognized standards now exist for the compression of video and text data. However, continued efforts are necessary in order to meet the particular requirements of the space science community. A NASA Workshop on Scientific Data Compression, held in 1988, recommended that it was of foremost importance to develop metrics of information quality and content for lossy compression schemes that would allow scientists to make intelligent choices regarding data compression vs. data loss. The same workshop also recommended that NASA continue the development of high-performance, high throughput flight-qualified data compression hardware that can be used on future missions. This latter recommendation was reiterated by the CODMAC in their 1988 report on Selected Issues in Space Science Data Management and Computation in endorsing data compression as an important component in an overall strategy addressing the management of high data rates and data volumes.

It is important to make astronomers familiar with the advantages and disadvantages of data compression, since they have not historically had to use it. Data compression techniques should be encouraged in NASA astrophysics flight projects, and prototyping of astronomy-specific techniques should be supported via the Software and Research Aids Program. Compression techniques should be considered for data within instruments, in temporary and permanent archives, and in transmission. The assumption is that astronomical instruments generate very high volumes of image data, and data transmissions from a Lunar outpost will most likely be extremely limited and/or schedule constrained. Thus there need to be

capabilities of decreasing the volume of data both within the instrument, when possible, and in data storage and forward.

The option of staged transmission should be considered, whereby low-volume "quick look" data is transmitted routinely, intermediate-volume data is transmitted periodically, and high-volume data is transmitted occasionally or on request.

The ramifications of archiving compressed data (either loss-less or with variable information content) on the query and retrieval process should also be addressed.

6.4.2 Automated Planning and Scheduling Tools

There is a need for proposal preparation aids and proposal management systems for missions. These must guide the users through the proposal process, provide information on expected performance and allow calculation of required input parameters for an observation. It would be desirable to see such generation aids integrated over missions so that there are common interface and basic functions. Similarly the management of the proposals, tracking their evaluation, notifying the proposers, and coordinating proposals could be broadly based. It should be possible to examine the observational program of several facilities to see what is being planned and to review what has already been done.

There are existing examples of "expert" planning and scheduling tools, including SPIKE, which was developed at STScI for long term HST planning and is being studied for EUVE and other missions; KDS, which is being studied for the ISO mission, and possibly the ROSAT system. Support should be given for studying these and other approaches and ascertaining how generalized they could become.

A Lunar Observatory will likely operate unattended for extended periods. Flexible, autonomous control systems are needed to ensure efficient utilization of scientific instruments and Observatory resources during these long periods. A central component of an autonomous control system is a planning and scheduling system for managing Observatory resources. The planning and scheduling of Lunar Observatory instrument and system resources is complicated by a number of factors which require automated systems for solution. These factors include complex and dynamically changing operational constraints due to scientific tasking, system health and status, and configuration of support resources. Each of these is discussed in turn in the following paragraphs.

Scientific tasking introduces several types of constraints on the planning and scheduling of Lunar Observatory resources. The first is oversubscription of system resources. Historically, requests for resources for science observations on space platforms have far oversubscribed system capacity. The process of scheduling space science activities in the past has primarily been done by hand and involved work-decades of effort even for relatively short observation periods such as planetary encounters. It is likely that high demand for space science resources from the science community will continue and be a severe constraint on the planning and scheduling of those systems. New automated tools are required which can assist Earth or Lunar-based scientists in managing initial resource allocation where requests vastly oversubscribe system capacity. Some computer-based tools have been developed to meet this type of need on other space projects. However, those tools are either obsolete, or tailored to specific missions or hardware. Existing tools (cf. the ST ScI SPIKE system for long-range planning) should be examined for possible generalization or adaptation. Autonomous scheduling systems are also required for managing those scarce resources at the Lunar Observatory to avoid overloading capacity when the system is forced to respond autonomously to dynamic situations, as described in the paragraphs below.

The second type of constraint introduced by scientific tasking is the need for a Lunar Observatory to respond quickly, efficiently, and autonomously to transient or emerging science opportunities and events. Once such an opportunity or event is detected and the desire to respond to it established, Lunar Observatory systems must be reconfigured dynamically, being careful to avoid oversubscription of system resources. Autonomous replanning which impacts the availability or configuration of resources must be accomplished with minimal disruption to existing schedules to avoid compromise of previously planned science activities. This calls for automated scheduling systems with a host of rescheduling strategies which can be chosen according to the response time available and other dynamic constraints of the situation at hand.

An autonomous planning and scheduling system is also required in order to respond quickly and efficiently to diminished capability due to failure or degradation of Lunar Observatory resources. An automated replanning system must reconfigure and reallocate Observatory resources in an attempt to recover and resume any activities which are disrupted by transient or permanent failures. An automated replanning system should also minimize the impact of reduced capability on future scheduled science activities.

Finally, the interaction of an unattended Lunar Observatory with external systems and support resources must be managed. This will require the planning and scheduling of data storage and data transmission facilities at the Observatory itself and throughout the data system on the Moon, on communications spacecraft, and on the Earth. Autonomous or semi-autonomous scheduling systems must be able to coordinate their actions to make efficient use of all resources.

6.4.3 Intelligent Tools to Assist in Data Analysis

Related to the need for autonomous operation and data compression is the desirability of developing intelligent systems to support the detection and analysis of interesting features in images, spectra, and temporal phenomena. One of the greatest challenges for automation in support of an unattended Lunar Observatory is the ability to recognize interesting and possibly transient science opportunities, and to respond to them through changes in observing strategy, configuration of Observatory resources, and coordination with other observatories. An additional challenge to Lunar science operations is to maximize the productivity of both resident Lunar Observatory scientists and scientists remotely operating or receiving data from the Observatory.

These challenges entail the ability to rapidly analyze the large volumes of scientific data which will be received from Observatory instruments. Currently, only a fraction of the data returned from space science missions is processed and analyzed in real or even near-real time (within days of acquisition). There are archives of science data from previous missions which have never been analyzed. Since the time of resident scientists will be a scarce resource and since the detection and analysis of interesting data may redirect observations or subsequent analyses, it is desirable that an automated science analysis system, as part of science operations, direct the scientists attention and effort towards "interesting" data and facilitate its interpretation.

In support of these requirements, an intelligent science analysis computing environment should be developed and should be tightly integrated with other Lunar science operations systems. Such an environment would include a variety of intelligent systems, data manipulation, and graphical visualization capabilities which serve as an "automated research assistant." They would facilitate rapid analysis of interesting science data, in the areas of preliminary review of data, suggested analysis methods, and cross-referenced information. The intelligent systems would incorporate both low-level pattern recognition algorithms for detection of interesting features in images, spectra, and temporal phenomena, and higher-level strategies and heuristics for performing preliminary analyses on interesting data. These systems must include abstraction and filtering capabilities to enable rapid evaluation of unusual data.

The results of automated detection and analysis of interesting data should be forwarded as "alerts" to an autonomous scheduling system in the Lunar Observatory for planning additional observations as well as to other scientific facilities or observatories for further evaluation and response.

These requirements can be addressed by recent developments in the fields of artificial intelligence, graphics, and data management. Some work has also been done in the astronomy community, in the areas of automated classification schemes and rule-based calibration procedures. To date, these technologies have not been integrated or widely applied to space science problems, and this must be a major focus of prototyping efforts. These considerations should not be taken as substituting for an individual researchers freedom to analyze data in unique and independent ways.

6.4.4 Instrument and Experiment Self-Monitoring Tools

Self-monitoring and diagnosis capabilities are absolutely necessary to ensure the continuous, reliable, safe, and productive function of an unattended Lunar Observatory. Automated self-monitoring systems for instruments and experiments should be developed which provide the primary functions of detection and characterization of faults, initiation of safeing actions, and communication of system health and status information to resident operations personnel and/or to external controllers in space or on the Earth.

An automated monitoring system is required in order to detect and respond to long and short-term trends in operational characteristics or parameters. The system should provide capabilities for automated troubleshooting throughout the end-to-end instrument, experiment, and support facility systems. An autonomous monitoring system must include accurate limit-checking, particularly when the limits are dynamic due to changing system configurations and loading. This capability in turn requires access to instrument simulations and specialized diagnostic analyses tailored to specific Lunar systems.

Information from autonomous monitoring systems should also be routed to automated scheduling systems at the Lunar Observatory to initiate replanning around lost or reduced capability in an instrument or subsystem when a backup system is not available. This illustrates also the overall requirement to systematically integrate a variety of intelligent systems in an encompassing automated Lunar Observatory operations control system.

Intelligent systems for telemetry monitoring and health analysis of multiple spacecraft subsystems have already been proven in space operations environments and are rapidly moving into mission operations as flight critical software in both the manned and unmanned planetary exploration programs. In most cases, these systems are very specialized for the particular monitoring applications. These systems should be generalized and extended to accommodate the requirements of automated monitoring of scientific instruments and experiments in the context of a Lunar Observatory operations control system.

6.4.5 Integrated Lunar Observatory Control System Testbed

The autonomous operation of a Lunar Observatory over long, unattended periods will require a variety of intelligent systems which must interact and coordinate their activities. To summarize: An automated planning and scheduling system is required which will manage Observatory resources and respond to dynamic, changing constraints issuing from the other automated systems as well as from external systems including humans. An automated data analysis system is required which will detect sudden, transient events of scientific interest and direct the automated planning and scheduling system with new tasking based on these scientific opportunities. An automated experiment and instrument self-monitoring system is required which will ensure the reliability of space operations systems by monitoring system health and status, and which will instruct the automated planning and scheduling system to reconfigure Observatory systems to avoid or recover from faults or degradation of capability.

A testbed is required for the purpose of integrating these and other required science operations systems for the Lunar Observatory. There is very little system design and engineering experience anywhere with control systems that include these types of automated systems. The development of an operations testbed will permit investigation of alternative autonomous operations systems designs, and will ensure autonomous subsystem interoperability and compatibility.

6.4.6 Standardized Instrument Command Structures

The existence of long-lived observatories on the moon will require a higher level of standardized instrument command structures than has been necessary in the past. A standard structure which is applicable to astrophysics instruments, and compatible with astronomical observatory operations will be necessary. This will require the development of systematic architectural models for the components of astrophysics information systems. Protocols must be defined for component interaction that reflects astrophysics domain-specific needs. The entire cycle, including mission planning, instrument control, instrument monitoring, and data analysis must be considered. The projection of these systems onto the computing domain, including the facets of communications, execution control, numerical processing, data management, and user interfaces must be taken into account in the development of these models and

protocols. It is essential that a framework be established in which standard and custom components can be mixed and matched to produce distributed, heterogeneous, evolvable systems.

